



Software Engineering Institute

Trusted Computing in Embedded Systems Workshop


Archie Andrews, CERT
Jonathan McCune, CyLab

March 2011

SPECIAL REPORT
CMU/SEI-2011-SR-002

CERT® Program
Unlimited distribution subject to the copyright.

<http://www.sei.cmu.edu>



This report was prepared for the

SEI Administrative Agent
ESC/XPK
5 Eglin Street
Hanscom AFB, MA 01731-2100

The ideas and findings in this report should not be construed as an official DoD position. It is published in the interest of scientific and technical information exchange.

This work is sponsored by the U.S. Department of Defense. The Software Engineering Institute is a federally funded research and development center sponsored by the U.S. Department of Defense.

Copyright 2011 Carnegie Mellon University.

NO WARRANTY

THIS CARNEGIE MELLON UNIVERSITY AND SOFTWARE ENGINEERING INSTITUTE MATERIAL IS FURNISHED ON AN "AS-IS" BASIS. CARNEGIE MELLON UNIVERSITY MAKES NO WARRANTIES OF ANY KIND, EITHER EXPRESSED OR IMPLIED, AS TO ANY MATTER INCLUDING, BUT NOT LIMITED TO, WARRANTY OF FITNESS FOR PURPOSE OR MERCHANTABILITY, EXCLUSIVITY, OR RESULTS OBTAINED FROM USE OF THE MATERIAL. CARNEGIE MELLON UNIVERSITY DOES NOT MAKE ANY WARRANTY OF ANY KIND WITH RESPECT TO FREEDOM FROM PATENT, TRADEMARK, OR COPYRIGHT INFRINGEMENT.

Use of any trademarks in this report is not intended in any way to infringe on the rights of the trademark holder.

Internal use. Permission to reproduce this document and to prepare derivative works from this document for internal use is granted, provided the copyright and "No Warranty" statements are included with all reproductions and derivative works.

External use. This document may be reproduced in its entirety, without modification, and freely distributed in written or electronic form without requesting formal permission. Permission is required for any other external and/or commercial use. Requests for permission should be directed to the Software Engineering Institute at permission@sei.cmu.edu.

This work was created in the performance of Federal Government Contract Number FA8721-05-C-0003 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center. The Government of the United States has a royalty-free government-purpose license to use, duplicate, or disclose the work, in whole or in part and in any manner, and to have or permit others to do so, for government purposes pursuant to the copyright license under the clause at 252.227-7013.

For information about SEI publications, please visit the library on the SEI website (www.sei.cmu.edu/library).

Table of Contents

Acknowledgments	iii
Executive Summary	v
Abstract	vii
1 Introduction	1
2 Workshop Goals	2
3 Workshop Format and Plenary Sessions	3
3.1 Summary of Plenary Session – Day 1	3
3.1.1 Introductory Presentations	3
3.1.2 Problem Panel	5
3.1.3 Break Out Session – Day 1	7
3.2 Summary of Plenary Session - Day 2	8
3.2.1 Current Research Reports	8
3.2.2 Case Study	10
3.2.3 Break Out Session – Day 2	10
4 Conclusion and Recommendations	12
4.1 Conclusions	12
4.2 Recommendations	12
Appendix A - TCES 2010 Agenda	16
Appendix B – Attendees	18
References	20

Acknowledgments

The authors of this report thank the members of the organizing committee for their valuable insight and recommendations. The committee members: David Fisher, CMU SEI, CERT, Virgil Gligor, CMU CyLab, Amy Kunkle, CMU SEI, CERT, Adrian Perrig, CMU CyLab, and Howard Lipson, CMU SEI, CERT, all gave generously of their energies. Thanks to all the participants in this workshop who gave us all their more valuable possession, their time, as they engaged in candid discussions and exchange of ideas both in forum and in numerous interactive opportunities. Special thanks goes to the excellent support we all received from the CERT and CyLab staff, especially Amy Kunkle and Samantha Stevick, who made the logistics for this workshop look easy. Finally, we owe special appreciation to Dr. Michael May, Associate Director for Software Technologies in the Office of the Secretary of Defense, for recognizing the challenges in this area and providing the funding and personal support to make this workshop of value to all.

Executive Summary

Despite progress during the past decades secure operation of computing systems and components continues to be a fundamental research challenge. Unfortunately, the increasing sophistication of defense mechanisms has resulted in a progressive evolution of increasingly sophisticated attacks. As a result, the majority of computing systems are still plagued by common vulnerabilities that can readily be exploited by attackers.

The concerted efforts of industry to resolve these challenges have resulted in the creation of new hardware security standards. Most visible is the Trusted Platform Module (TPM), specified by the Trusted Computing Group (TCG), a not-for-profit standards organization. The TCG effort has been so successful that hundreds of millions of current computing platforms today are equipped with TPM chips.

The Trusted Computing in Embedded Systems Workshop convened to begin discussion on the necessary research to advance the application of trusted computing technologies in embedded systems. This workshop provided the opportunity to focus on embedded systems while incorporating the relevance of other related disciplines to foster collaborative efforts in this area.

Presentations and discussion addressed the capabilities and limitations of effectively employing trusted hardware-enabled components in embedded systems. This included, but was not restricted to, the following areas:

- new research and development in enabling trust in embedded systems
- methods and techniques for establishing trust in embedded systems
- lessons learned from research and development projects on embedded systems security
- gaps in current research

Based on the workshop presentations and candid exchange of ideas and concerns, the authors of this report developed recommendations for potential research directions for improving trust in embedded systems.

- Research is needed on isolation and memory management methods to improve the level of trust feasible in 8-bit microcontrol systems today.
- Coherent trust models based on sound criteria for, and principles of, trust are necessary to support both acquirers and researchers.
- A reference implementation of end-to-end use of trusted computing in an embedded system would be a valuable community resource to explore the characteristics of such a system and to understand necessary compromises, shortfalls and limitations.
- Tool support is necessary to aid embedded development and design teams in their consideration and incorporation of security and trust into their requirements, specifications, designs, and implementations.

- Building a community that spans the safety, security, dependability, and trust communities requires research to identify the points of intersection and diversion as well as encouraging sponsorship and leadership.
- Research is needed to appreciate the relationship and dependencies inherent in cyber-physical systems and to evaluate how that interrelatedness impacts trust.

Abstract

This report describes the November 2010 Trusted Computing in Embedded Systems Workshop held at Carnegie Mellon University. This workshop brought together various groups concerned with advancing research into improving the trustworthiness in embedded systems. The workshop format provided the opportunity to focus on embedded systems while examining the application of related trust technologies in order to foster collaborative approaches and information exchange in this area. Presentations and discussion addressed the capabilities and limitations of effectively employing trusted hardware-enabled components in embedded systems. This included, but was not restricted to, the following areas: new research and development in enabling trust in embedded systems, methods and techniques for establishing trust in embedded systems, lessons learned from research and development projects on embedded systems security, and gaps in current research. The workshop resulted in identification of gaps in current research and recommendations for potential research directions.

1 Introduction

The TCES workshop, held on November 9 -10, 2010, in Pittsburgh, PA on Carnegie Mellon University's campus convened to explore the practical application of trusted computing technology to embedded systems and to inspire collaborative opportunities and future information exchange among the government, industry and academic participants. The workshop was sponsored by the Office of the Secretary of Defense, Director of Defense Research and Engineering (ODDR&E). CERT and CyLab at Carnegie Mellon University jointly hosted the meeting. The workshop was targeted at developing a shared appreciation for the application of trusted computing to embedded systems to include currently available means of establishing trustworthy embedded systems, technology gaps in trustworthy computing for embedded devices and systems, and promising areas for future research and practices. A workshop steering committee from CERT and CyLab agreed on the abstract and specific goals for the workshop. The workshop agenda focused on presentations and research papers that help characterize the application of the current trusted technologies in embedded systems and frame the necessary research to advance the field. Reports presented in plenary session addressed both research challenges and the current state of research. The breakout session discussion aimed to identify the difficult challenges in supporting trusted computing in embedded systems.

Individuals attended the workshop from government, industry and academia. Participants were invited based on their active engagement in research and/or acquisition of trusted systems or embedded systems. (See Appendix B, "Workshop Attendees.") This disparate group of participants shared their perspectives on the current needs, capabilities, and limitations of applying trusted computing technologies to embedded systems that operate in increasingly hostile environments.

2 Workshop Goals

While individual technologies already provide enhanced security capabilities, the greatest challenges still lie ahead in combining these technologies in ways that will enable secure, robust, and manageable trusted systems and components to become the norm for the next-generation of embedded-systems. The key to addressing these challenges appears to be research in trust-enabling infrastructure and application to embedded systems.

This workshop focused on identifying, supporting, and driving interdisciplinary research activities to address the security challenges relevant to embedded systems. The two-day workshop was intended to foster further collaborative research among academia, industry, and government technologists, linking industry's experience in the development of Trusted Computing technology, academia and CERT's IT security skills with the government's desire to enhance the security of the national information infrastructure.

Goals for the workshop included:

- Bring industry (producers and system integrators) and government practitioners (consumers and users) in embedded systems together with researchers in trusted computing to explore feasible and practical application of trusted computing technology to embedded systems.
- Build a shared understanding of the following:
 - currently available means of instilling trust in embedded systems,
 - practical limitations bounding that trust within current implementations,
 - promising areas for future research and practice.
- Identify research areas necessary to address the security challenges relevant to embedded systems.
- Produce a summary report on promising research areas, benefits to DoD, and projected timeframes (e.g., for practical application of the proposed research).
- Make contacts for future information exchange and possible collaborative research projects.

3 Workshop Format and Plenary Sessions

The TCES workshop agenda (see Appendix A. “TCES 2010 Agenda”) was a combination of plenary sessions and breakout sessions. The sessions were designed to develop a common perspective on currently available means of instilling trust in embedded systems, practical limitations bounding that trust within current implementations, and promising areas for future research and practice. This workshop brought together disparate groups to gain a broad perspective of the problem. Each workshop participant was invited based on their active engagement in research and/or acquisition of either embedded systems or trusted systems.

Day one focused on practice with keynote and panel sessions covering challenges and current problems to lay a foundation for shared understanding. There were sessions on cyber-physical systems (as an example of networked embedded systems) and a background session on the promise of trusted computing and recent advances in that area. Five panelists participated in a problem panel that laid out the problem-set from the view of the industry producer and integrator, the government user, and the researcher. The breakout sessions on day one developed plausible attack/defense scenarios directed toward embedded systems to explore needed trust capabilities.

Day two focused on research with reports given by workshop participants engaged in ongoing research on their current work as well as a case study on a directed attack, the Stuxnet worm. The case study on Stuxnet was a late addition to the agenda to ground the discussion on potential future directions in the reality of an ongoing investigation of a real and recent attack against an embedded system controller. The second day concluded with all participants engaging in a brainstorming session to identify current gaps in research in the area.

3.1 Summary of Plenary Session – Day 1

The initial plenary session framed the challenges and potential gains for necessary research.

Michael May, Associate Director for Software Technologies in ODDR&E, welcomed the participants and stressed the potential value to the DoD of finding answers to the hard problems of building trustworthy embedded systems. He pointed out that his role was to act as an advocate and connecting agent for required research with the appropriate agencies within each of the DoD branches of service.

3.1.1 Introductory Presentations

Introductory presentations were used to set the context of the areas of embedded systems and trusted computing in preparation for discussions of the possible overlap and future opportunities.

Archie Andrews, SEI’s CERT, presented the challenges that need to be addressed in order to build trustworthy embedded systems. Andrews proposed that there are underlying problems awaiting solutions such as operating within a known malicious environment, maintaining trust in environments with autonomous components, and cooperative detection and mitigation of potential failures. He described the opportunities from the workshop of embedded systems researchers gaining an understanding of the relevance of recent advances in trusted computing, trusted computing researchers appreciating the characteristics and limitations of embedded systems, and members from both areas collaborating to explore how to build trustworthy embedded systems.

Al Mok from the University of Texas at Austin gave a presentation on Trust in Cyber-Physical Systems. Mok pointed out that Cyber-Physical Systems (CPS) predominantly rely on embedded system components but by

definition will include the tight conjoining of and coordination between computational and physical resources with the attendant complexity of such a “mash-up” and the implications on establishing trust in such a system. His presentation revealed the many facets of trust inherent in an event chain that involves hardware/software interactions. He explained that the challenge to trustworthiness extends beyond security (real time interaction) to include reliability (the failure semantics challenge) and interoperability (the open system environment challenge). He spoke to the complexity of failure semantics (how a system behaves when a failure occurs) introduced from the interrelatedness of cyber and physical components in which the failure of one component, either cyber or physical, can affect the other in complex ways. Mok asserted that specifying failure semantics that are both strong and implementable is a hard problem. In reference to the interoperability inherent in the mash-up of cyber and physical in CPS, Mok pointed out that “Open Systems” carries multiple meanings. He explained that open systems, when used in the context of computing, implies some degree of interoperability, portability, and open systems standards. He contrasted that idea with open system, when used in systems theory, refers to a system where matter or energy can flow into and/or out of the system as opposed to a closed system in which energy can enter or leave but matter may not. Both meanings are applicable to CPS. Mok pointed out that this open system environment has implications in both policy and implementation issues and these issues may affect the trust and certainty of the system. Mok also spoke to the security challenges and potential opportunities for defense in CPS. He explained that as CPS systems grow more complex and interactive with the environment then there will be more vulnerabilities, more failure modes, more overt channels and more unexpected interactions. At the same time, the unalterable laws of physics determine the dynamics of the physical system; this can be in the defenders’ advantage. Mok asserted that there may be opportunity for damage containment in CPS by defining new failure semantics that incorporate CPS domain-specific semantics based on the inherent physics of the system under control and the temporal dimension (e.g., time-to-detect, time to correct action). Mok’s presentation briefly introduced the pros and cons of trustworthiness of controls in a wireless and a wired CPS as an example of some of the considerations for understanding the nature of trust and trustworthy in the context of CPS and embedded systems.

Jonathan McCune, CMU’s CyLab, gave a presentation on Trusted Computing Concepts and Recent Advances to introduce the trusted computing landscape to the participants for background and to add trusted computing techniques for consideration in new embedded systems. McCune’s central messages were that trusted computing mechanisms enable fundamentally new properties and trusted computing mechanisms provide new primitives to build secure systems. He motivated these central messages by posing a controlling question: how does one get a low-level guarantee that everything is as intended on any remote device, in other words, how can one remotely verify trustworthy device operation. He then laid out an axiom – every system has at least one more flow. Additionally, he presented some assumptions, such as the concept that a remote adversary can launch networked-based attacks, hardware under control is trusted (i.e., physical anti-tamper mechanisms are complementary but outside scope). Given those considerations, McCune posed the question - how can we establish trustworthy device operations. The question was especially pertinent to the workshop because remote attacks constitute the most serious new threat to networked embedded systems such as those described by Al Mok. To familiarize the workshop with current approaches, McCune examined the advantages and disadvantages of a number of different approaches such as program code in ROM, secure boot and virtual machine isolation. That was the springboard to a discussion of the advantages and applications of remote attestation based on the Trusted Computing Group’s Trusted Platform Module (TPM). He also introduced the basic TPM functions and their potential application in a number of environments of interest to workshop participants. Some applications may be hardware based attestation of Linux static configurations and isolated execution environments, and software based attestation.

3.1.2 Problem Panel

Following the introductory presentations, a problem panel session was held to expose the problem from the perspective of the industry producers and integrators, the government user, and the researcher. The panel consisted of Helen Gill, co-chair of NITRD High Confidence Software and Systems Coordinating Group, National Science Foundation; Robert Thibadeau, Senior Vice-President and Chief Scientist, Wave Systems Corporation; Ron Perez, Fellow and Security Architect, AMD Corporation; Tom Shaffnit, Ford Motor Company; and Duane Gilmour, Chief Computing Architectures Branch, Air Force Research Laboratory. Adrian Perrig of CMU's CyLab introduced the panel and invited each member to talk about the problems that kept them awake at night.

Helen Gill used the backdrop of cyber-physical systems to introduce the challenges that span the range of embedded systems. She began by defining cyber-physical systems (CPS) to include aspects of computing embedded into every physical components, possibly even into the materials, networked at every scale, with that computational core as an embedded system usually demanding real-time response. She explained that while CPS is not traditional, post-hoc embedded/real-time systems, it will display cyber capability in every physical component, complexity at multiple temporal and spatial scales, dynamic reorganization and reconfiguring, as well as unconventional computational and physical substrates - all this while requiring operations to be dependable and certified in some cases. Gill challenged the workshop participants to look beyond current embedded systems to the next generation, for example expanding the scope (simple to complex, hardware-software to full system considerations). She also stressed the increased necessity for end-to-end assurance in interoperable and open systems. Her conjecture is that dealing with the end-to-end problems will require integration of previously separate IT research areas such as real-time embedded systems, control theory, networking and communications, security and privacy, and human-computer interaction. Gill explained that the trusted computing in embedded system challenges stem from the following problems:

- Open, complex dynamical systems;
- coexisting and joint accreditation of mechanisms such as security, real-time task management, control, networking and autonomy;
- understanding the perspective of process and behavioral dynamics;
- appreciating authority management with respect to control authority; and
- widening / opening the trusted computing research community

Robert Thibadeau related that each year he searches for a magic bullet, i.e., a principle that would cause sweeping change for the better. He cited the US Constitution and the Ten Commandments as examples. For this fiscal year he proposed two such principles related to trusted computing and embedded systems, namely:

- Every device should be able to uniquely sign and have a corresponding X.509 public key certificate signed by the manufacturer. Thus each important agreement will be based on proof that both parties have agreed in a way that neither party can deny or repudiate.
- Common criteria level 1 is self certified. The implication of this principle is that neither a lab nor an external process or approval is needed for the most basic set of security assurance requirements. Thibadeau's belief is that removing the obstacles and expense to considering security assurance requirements is a strong incentive to draw companies toward the right path. He also proposed that laws would have to be passed that make fraudulent self-certification punishable.

Ron Perez provided an industry perspective on the challenges facing trusted computing technologies and embedded systems. He began by recognizing that the embedded systems considerations are different than the general purpose computing model that the Trusted Computing Group is focused on. He pointed out that considerations such as cost sensitivity (even a \$1 added to the BOM maybe too much), security requirements (what is the ROI on security investment), and usability (opt-in is probably impractical, requiring on by default and transparent usage) are important and different considerations for the embedded systems community. Another difference he observed was the different ecosystems. For example, general purpose systems require end-to-end support for software and services, whereas embedded systems are generally closed ecosystems with standards driving and supporting composability. Another example was the feasibility of attestation, where the general purpose platform is highly dynamic both at load-time and at runtime while an embedded system has a more static workload on purpose-specific platforms. After identifying those distinguishing characteristics, Perez went on to enumerate the challenges unique to embedded systems that bear on imbuing trust in embedded systems. The challenges he listed include environment and reliability requirements, form factor constraints, diversity in devices, hardware attack mitigation and real-time requirements, as well as debugging, failure and testability issues. Perez noted that it is too easy NOT to use trusted computing technologies in the embedded systems world, as a result of cultural biases such as “security through obscurity” and rolling one’s own solutions.

Tom Shaffnit spoke to the challenges of designing a system that relies on cooperating embedded devices with integrated security for reliable and predictable performance. He used as an operational example the vehicle-to-vehicle (V2V) interoperability project involving most of the major automotive manufacturers. The V2V project combines vehicle communications with GPS to provide a new safety sensing capability. According to Shaffnit, the challenge is to design for interoperable security features amongst all automotive manufacturers that eventually extend to serve as a vehicle-to-infrastructure safety system. Shaffnit explored the complex interrelatedness of trust and embedded systems by noting that the V2V project system design requires an understanding and appreciation for the questions of policy issues (e.g., control, rights, privileges, privacy, and security) as well as technological challenges.

Duane Gilmour gave a service perspective on trusted computing concerns. While Gilmour specifically addressed the concerns of the Air Force, his concerns appear to be equally applicable to the other services. Gilmour expressed what keeps him up at night as a series of statements and unresolved questions. Those concerns include:

- establishing a root of trust
- dealing with supply chain issues, specifically off-shore manufacturing
- ramifications of adding security to off-the-shelf technologies versus requiring security to be built into new technology
- appropriately evaluating hardware and software trust as well as answering the question – can we achieve trust with untrusted components?
- protecting data, given that our current supply chains must rely on off-shore suppliers and knowing that we will process critical data on untrusted hardware.
- defining security implications given that we are moving to multicore technologies.
- determining if there are potential unanticipated complex interactions of embedded software on multicore technologies?

- determining if we can develop provably correct code and compose trusted complex systems?
- determining if we can formally verify complex systems for trust, since test and evaluation can never be exhaustive.
- determining how we can ensure that we can fight through a cyber attack while completing a mission given that attacks to our national defense infrastructure will occur.
- identifying that a system has been compromised.
- determining how much trust is enough when considering the tradeoff decisions between trust and cost in overhead and application dependencies.
- determining if we can achieve trust in autonomous systems?

The comprehensive nature of the concerns and issues that the problem panel raised made sure that no one slept well that night.

3.1.3 Break Out Session – Day 1

Following the problem panel the workshop participants broke into three separate groups with each group challenged to describe a realistic scenario that relied on embedded systems and required security. Each group was free to capitalize on the expertise and interests of their constituents to develop a scenario. Once the scenario was developed, the group put on a black-hat to describe how the particular system described could be attacked. After exhausting potential attack scenarios, the group became white-hat defenders to describe how to protect the system described in the initial scenario. The teams were comprised of equal distributions of researchers, acquirer/users, and producers. The teams received no guidance or restrictions on what scenario to describe, making them rely on advocacy to emerge from the expertise within each of the teams.

The intention of this exercise was not to derive realistic threats or defense strategies. The goal was to encourage the participants consider the range of the threat environment to embedded systems and the hard challenges of securing those systems.

Two teams chose to describe a vehicle system as their system to first attack and then defend this system. The third team chose to describe a remote health care system. Every team had similar concerns with the threats and risks associated with respect to the dissimilar scenarios. All of the teams were consistent in their concern that increased connectivity and reliance on embedded systems, heightened the risks of either intentional or unintentional disruptive events. All of the teams took both a narrow and a broad perspective of the objectives of a disruptive event and potential consequences. Both of the automotive groups listed harm to national transportation infrastructure and economic disincentive to manufacturer, as well as local events such as blocking roads and tunnels. The health care group listed events such as reputation of institution, disruption of caregiver routine, and reputation of device manufacturer. All of the groups mentioned issues with surety of both hardware and software supply chains, fragileness of systems and the potential engineered impact to corporate reputations. The groups were easily able to describe the potential for harm to individuals such as malicious control of safety devices – airbags, brakes, and medication delivery devices; and harm to corporate concerns – manufacturers’ reputation, malware injection during routine maintenance, violation of patient privacy to damage institution. As for defensive measures, all groups addressed “reasonable test”, which is the concept that embedded devices should abide by the expected laws of physics and demonstrate behavior within expected norms. Operating beyond those norms – too high a dosage, shifting into reverse at 60 mph – should be identifiable by checks and balances within the embedded systems. The health care group included exploring opportunities for increasing the surety of the delivery devices to mitigate human errors. One of the automotive groups brought

up the issue of potential harm to national infrastructure and remarked that there were voids in responsibility and technical acumen to recognize, plan for, and mitigate disruptive events at the infrastructure level. The post-breakout observation was that it was remarkable how the problems and the potential solutions were so similar, even for such apparently dissimilar applications.

3.2 Summary of Plenary Session - Day 2

Day two of the workshop focused on research. Participants were invited to inform the workshop attendees of their current research and early results.

Philip Koopman, CMU, requested time for an impromptu session to build on the practice-oriented presentations of day one. Koopman wanted to relate current embedded systems to security and trusted computing. He referred to this as a reality check on embedded systems constraints. Koopman provided the workshop with a sense of the densities of the various types of CPUs in use worldwide. He stressed that designers and integrators optimize embedded systems for low total system cost. A TPM-type add-on would double the cost of most 8-bit systems. Koopman went on to give examples of potential and real intersections between innocuous embedded systems, such as those used for entertainment, and embedded systems with critical control functions. As an example, Koopman discussed how the infotainment system on an airplane and the in-flight control systems are separated only by a potentially fallible gateway. He emphasized the fact that this threat is real and should be of high concern. He compared security to dependability and trust pointing out that in most cases faults are assumed random and independent for hardware, the environment is assumed to behave as expected, and software is assumed to perform without fault. He acknowledged that all of these factors affect classic dependability. His conjecture is that security can be viewed as classic dependability with the added complication that attackers can intentionally correlate induced faults and attackers can manipulate environments to produce exceptions. Koopman's overview set a solid foundation for continued discussion of relevant research in both embedded systems and trustworthy computing that filled the rest of the day.

3.2.1 Current Research Reports

Eleven researchers were prepared to share their current research with the workshop attendees. The following is a very brief abstraction of the research reports. The reader should refer to the speakers' presentations found on www.cert.org/tces (see Appendix A for links to the presentations).

Tom Eisenbarth, Florida Atlantic University, presented his work on embedded security. Eisenbarth's presentation focused on the challenges of implementing crypto in embedded devices considering limitations in computing power, energy and memory as well as protecting against side channel attacks. Based on his research, Eisenbarth concludes that the adversarial models are different and porting trusted computing concepts needs to be done with care. The new directions of his research are leakage aware protocols and crypto schemes suited for embedded systems.

Kevin Fu, University of Massachusetts, presented ongoing research to understand the challenges of building trustworthy embedded systems, in particular trustworthy medical devices. Fu spoke about the unique challenges of embedded medical devices and the work that his group is doing to understand those challenges.

Rick Han, University of Colorado at Boulder, introduced his group's research on privacy in context aware (mobile sensor) social networks. His current research is focused on addressing the risk to preserving the privacy of a participant in a complex community using multiple assets (picture a social network user on Facebook, sensor networks and mobile networks) by building an anonymity layer.

Trent Jaeger, Penn State University, spoke to research at Penn State University on high integrity computing for embedded systems. The goals of that group's research is to understand how to specify integrity, how to deploy systems that meet the specifications, and how to maintain system integrity. His group is currently examining information flow models to understand how to establish and maintain integrity-relevant operations. The group is also investigating how to apply integrity concerns to distributed embedded systems.

Ramesh Karri, Polytechnic Institute of New York University, spoke about his group's research on securing the core root of trust. His group is challenging the assumption that the core root of trust embedded in hardware is secure. He pointed out that chip foundries and all aspects of IC design are increasingly migrating outside the US and subject to reverse engineering, Trojans, cloning and counterfeiting.

Philip Koopman, CMU, presented what he labeled a quick tour of current research in secure and safe embedded systems at the Electrical and Computer Engineering department at Carnegie Mellon University. His group's research is concerned with issues such as security and safety concerns with integrating gateways separating embedded systems with enterprise systems, authenticating embedded network messages, and potential impacts of embedded virtualization.

Farinaz Koushanfar, Rice University, presented her research on the physical dimensions of embedded security. Her group at Rice is investigating the use of physical one-way functions and physical unclonable functions to enhance the security of embedded devices.

Insup Lee, University of Pennsylvania, discussed his group's research on medical device cyber physical systems (MCPS). Lee pointed out that medical devices are being networked for integration and interoperability to increase functionality. As a result, they are evolving to become systems of systems and inheriting all the complexity and vulnerability of that domain. He enumerated four areas of concern for security attacks on MCPS: patient's physical security, patient's data security and privacy, the physical security of the device and the institution's data security and privacy.

Captain Roy Porter, Air Force Research Laboratory, gave an overview of the information directorate of AFRL and their research interests in trusted computing and embedded systems. Captain Porter made the point that the Air Force requires systems that are resilient against attacks and will be able to "fight through" to mission completion. He explained that to accomplish this requires high-assurance systems for embedded applications along with secure encapsulation of subsystem components. In his presentation slides, Captain Porter identified specific targets for trusted computing technologies in the areas of trusted software and trusted hardware.

Ed Suh, Cornell University, introduced the work that his group at Cornell is doing on hardware support for trustworthy embedded systems. Suh pointed out that the context of networked embedded systems, with many devices collaborating, requires a means of establishing and maintaining trust while recognizing that individual components will get compromised. He asserted that in today's embedded systems network, software and physical attacks are all possible and "fail-stop" may not be an option. In addition, according to Suh, embedded systems have constraints on real-time performance and energy consumption. Some of the areas researchers at Cornell are investigating to address these challenges include:

- establishing trust among devices by identifying the hardware via device attestation (physical unclonable functions for unique characteristics of each IC),
- providing a secure execution environment by using secure processors (single-chip processors with self-authentication) to support physical tamper resistance,

- recovering on failure through early detection via fine-grained monitoring (Dynamic Information Flow Tracking)
- fast roll-back to a trusted state via hardware check pointing (hybrid memory for instant check pointing and roll-back).

Doug Blough, Georgia Tech, presented a technical overview of MedVault. MedVault ensures the security and privacy of electronic medical records. While not directly applicable to embedded systems, it was clear from Doug's explanation on his group's work on protecting patient records such as dynamic policy combination, selective disclosure with source verifiability, and redactable signatures, could have bearing on areas of TCES research.

3.2.2 Case Study

We invited Liam O Murchu, Operations Manager for Symantec Security Response, to present a case study on the work that Symantec has done in analyzing Stuxnet. The case study was requested as a special, unscheduled presentation to this workshop because of Stuxnet's methods of distribution and propagation, its effect on industrial control systems, and the relevance to being able to build and maintain trusted embedded systems in a malicious environment. According to O Murchu, Stuxnet is the first known cyber attack incident that exhibits this degree of sophistication in dissemination, distribution, and precision targeting. O Murchu explained that the Stuxnet worm spreads indiscriminately but the malware payload is specifically targeted. O Murchu made clear that the attack was directed against programmable logic controllers in industrial control systems but the infection media was spread via general purpose computing networks supposedly isolated from the target. The analysis and understanding of Stuxnet continues and the reader should refer to more up to date sources for detailed explanations of this worm. The relevance to this workshop is the demonstration that a sophisticated targeted attack can reach its intended destination in an embedded system by exploiting emergent non-deterministic effects.

3.2.3 Break Out Session – Day 2

The initial plan for this break out session was to capitalize on the information we all received from the research reports and break into separate groups to discuss three subjects:

- Gap Analysis - what are the problem areas that the research is not currently addressing
- Baseline Capabilities - what are the fundamental characteristics that should be possessed by any embedded component claiming to consider trustworthiness
- Barriers - what are the barriers to adoption

However, due to the unplanned but welcome case study on Stuxnet, the group decided to stay in plenary session and use the time remaining on day two to have a moderated group discussion around the topic of gap analysis.

The following is a summary of the key points of the group's discussion related to identifying problem areas in leveraging trusted computing in embedded systems.

- A reference implementation of end-to-end usage of trusted computing in an embedded system would serve as a useful proof of concept.
- The questions of usability versus operational overhead that are engendered by trusted computing, especially in an embedded system environment, are currently unanswered.

- Requirement specifications for trusted computing in embedded systems need better definitions and better understanding by the acquirers.
- Tools need to be developed to assess hardware, especially if it is made offshore. Some potential tool candidates are boundary scanning tools able to look for extra gates. This idea may be closely associated with the work on IC fingerprinting from IBM Research [Agrawal 2007].
- Embedded development and design teams seldom have designers with computer science degrees and rarely have security experts. Research is needed to identify tools to aid the design and development teams such as CAD tools, checklists, and other design tools and processes.
- Trust must go beyond security and correctness otherwise the next generations will be condemned to brittle solutions. Coherent trust models that relate to identified threat models are missing. Research should develop the underlying principles rather than relying just on anecdotal scary stories. There needs to be a quantitative notion of trust (recognizing that finding lower bounds is hard in computer science). While it is easy to find witnesses to lack of trust, it is hard to find evidence of rewards of building reliable trusted systems.
- Risk analysis measures and techniques are needed that tell us when to say no.
- At present, the safety and security communities rarely talk. That may be a cultural issue, but we need to find ways to remove the barriers and encourage coordination and collaboration.
- Research is needed to develop model-based checking for development and enforcement of trusted attributes in embedded systems. Does there exist a set of “synthetic” natural laws that apply to trusted computing in embedded systems?
- Research is needed to determine reasonable inputs and outputs of a trusted embedded system. These may be somewhat analogous to type safety but research should test that claim. Investigation of abstractions and languages is needed to cross the barriers from the real-time bounded world of embedded systems to the transactional world of general purpose computing.
- Because reality requires a bounded amount of money in real products, there should be research on intelligent ways to make tradeoffs to get the highest practical security as well as some type of measurement of return on investment. Given the dominance of 8-bit microcontrollers in the embedded systems world, there needs to be research on what security enhancements are appropriate. Supporting research in memory management restrictions at the 8-bit level will support understanding possible and practical enhancements. Understanding isolation mechanisms applicable to the 8-bit microcontrollers will be useful.
- Simple tasks are hard to do today. Answering questions such as, “is my code still there on a remote device and has it been modified,” is hard today. Research is needed on how to answer such seemingly simple questions given the characteristics of embedded distributed systems.
- Research is needed to develop theoretical constructs for IT and physical systems operating together. Discovering a way to make IT and physical complementary rather than exacerbating would be a significant breakthrough.
- The workshop participants challenged themselves to identify some challenge problems for the group to look at.

4 Conclusion and Recommendations

4.1 Conclusions

This workshop met its stated goals. The goal of the workshop was to bring industry, government and researchers from the communities of trusted computing and embedded systems together for a candid exchange of ideas and problems. The presentations and ongoing discussions were intended to build a shared understanding of the means of instilling trust in embedded systems, the limitations that currently bind that trust, and promising areas for future research. The candid discussions in the plenary, the breakout sessions, and in the extemporaneous sessions, were meant to identify relevant research areas. The presentations and discussions during the workshop were the primary ingredients that lead to the creation of this summary report and the authors' interpretation of promising research areas identified, as well as projected timeframes for practical application of the research areas. The workshop provided the venue for researchers, producers and users to make connections for future information exchange and possible collaborative projects.

The theme of the workshop and the challenge of the field attracted many of the thought leaders from industry and government working to instill trust into embedded systems and many of the researchers exploring new approaches to potential solutions. The aim of the workshop organizers was to bring in more government participants to describe the challenges they face each day in system specification, acquisition, and implementation. Unfortunately, due to the timing of the workshop soon after the start of the new government fiscal year, travel for government participants was difficult. The lesson we learned from this is to schedule future such workshops in the February – July timeframe to maximize the opportunity for government participation.

The primary reason for inviting this particular set of participants was to build a shared understanding of the embedded systems, trusted computing technology, and the hard challenges these two communities of interests face. The workshop accomplished this goal not only via the presentations of the participants related to their areas of expertise and concern, but also the free-flowing dialogs and contributions by all. Each of the invitees was willing to share areas of their work as well as explore promising areas to address current limitations. A benefit of this workshop was that by assembling the differing communities of interest into a focused workshop they were able to explore promising areas for future research. Even if total shared understanding was not achieved, a shared appreciation of the hard problems facing the community was certainly accomplished.

Against the background discussion of the current state in the fields of embedded systems and trusted computing, the current research efforts and directions, and the example of a directed attack were all contributors to a targeted discussion of necessary research areas. Focused attention, especially through collaboration of researchers and individuals concerned with practical application of the research findings, could result in significant progress on this set of challenging problems.

4.2 Recommendations

At the risk of omitting some important ideas, the compilers of this report feel that the two days discussion and the 15 key points identified by the group on day two distill down to six recommendations for attention. Without ascribing any priority based on ordering of the list, the recommended areas of focus, some background on the authors' reasoning in deriving these recommendations from the workshop's ideas, and the authors' estimation of a probable timeframe for reaching useful results follows:

- Recognizing that 8-bit microcontrollers dominate the embedded systems of today, research is needed on what isolation and memory management methods may be applied to address the level of trust feasible in such systems today. It was clear from Philip Koopman's presentation that the preponderance of embedded systems are simple, 8-bit devices with all the incumbent limitations on space, weight, cost, and power. This class of devices has traditionally assumed either a benign environment or the absence of an intentional adversary. During the initial breakout session it became clear that inventive experts could develop reasonable attack scenarios directed against this type of device. The authors felt that the density of devices provides the opportunity and, as explored briefly in the scenario-sessions, the motivation will be found. A target-rich environment will not go un-targeted. The effort to define existing vulnerabilities and potential cost-effective remedies is near to mid-term research. Success in describing the problem and providing sufficient insight into possible solutions has the potential to influence the chip design and the fabrication market.
- Coherent trust models based on sound criteria for, and principles of, trust are necessary to support both acquirers and researchers. Defining principles of trust provides a guide against which to measure specifications, requirements, and results. The need for developing sound principles was raised during the gap-analysis discussion and fostered by numerous side-bar discussions during the networking sessions. The desire is to have a precise set of criteria that will allow measurement of success or failure in meeting those criteria. If the criteria expressed as a principle, cannot be met, then the community needs to amend their composite notion of achievable trust. The authors feel that the initial research in developing a set of principles of trust should be a near-term objective. The creation of such a set will hopefully drive a longer term discussion, debate, and experimental implementations that will in turn test and refine the foundational principles.
- A reference implementation of end-to-end use of trusted computing in an embedded system as a proof-of-concept would be a valuable community resource to explore the necessary characteristics of such a system and to understand necessary compromises, shortfalls and limitations. This suggestion arose during the gap-analysis brainstorming session. The opportunity to template an embedded system as a proof of concept demonstration that could be emulated and improved by others resonated with the authors of this report. The model could be similar to the capability that Kevin Fu, University of Massachusetts, described in the domain of trustworthy medical devices. This research could be initiated by a number of groups in the near-term with the aim of demonstrating numerous specialized applications that are applicable to broad classes of problems.
- Tool support to aid embedded development and design teams in their consideration and incorporation of security and trust into their requirements, specifications, designs, and implementations. Embedded systems design and acquisition teams are rarely staffed with needed expertise. Tools and processes that aid or contribute to adoption of security in embedded systems are needed. Recognizing that few teams have security experts or computer scientists, there is an opportunity to bring security best practices to the embedded design community. Tools and processes identified by this group range from hardware validation (is trust in off-shored hardware justified), usability (operational overhead accompanying increased levels of trust), risk analysis measures and techniques, and engineering tradeoff decision criteria. The need for research and ultimately solutions in this area came up during the gap analysis session. This long-term research requires evolving the tools as the required attributes of trusted systems get defined and the technology emerges to support them.

- Building a community that spans the safety, security, dependability, and trust communities requires research to identify the points of intersection and diversion as well as encouraging sponsorship and leadership. One of drivers that prompted this workshop was the conjecture that the community of embedded systems researchers and practitioners and the community concerned with trusted computing rarely exchanged ideas or shared perspectives. During the discussions, both planned and unplanned, the observation was made that the same siloed mentality applies to a number of fields that would benefit from sharing knowledge and experience. Defining the covalent (where electrons are shared) aspects of the desirable conditions should be near-term research involving collaboration between members from the trusted computing and the embedded systems community to build on the previous work of [Avizienis 2004] and others. Establishing and maintaining a dialog among these communities to address such issues as abstractions and languages that can cross between transactional and embedded computing will require workshops and venues for shared experiences and interactions.
- The interrelatedness of IT and physical systems has direct relevance to the DoD and other government entities (e.g., DoE, DHS, DHHS, etc.). Research is needed to appreciate the relationship and dependencies inherent in cyber-physical systems and to evaluate how that interrelatedness impacts trust. As was pointed out by numerous speakers (e.g., Helen Gill, Al Mok, and Insup Lee), this is long-term research dealing with a wide-spectrum of issues spanning the fields of embedded systems, system of systems, trust enablers, and failure mitigation. Although NSF is funding academic research in this area, the direct application to DoD systems such as unmanned and autonomous system, makes this area of high importance to the DoD.

Appendix A - TCES 2010 Agenda

URLs are valid as of the date of publication.

Tuesday, November 9, 2010

Focus is on Practice

Welcome, Michael May, OSD

Hard challenges, Archie Andrews, CMU CERT
<http://www.cert.org/tces/pdf/archie%20andrews.pdf>

Networked embedded systems, Aloysius Mok, University of Texas-Austin
http://www.cert.org/tces/pdf/al_mok.pdf

Trusted computing concepts and recent advances, Jonathan McCune, CMU CyLab
<http://www.cert.org/tces/pdf/jon%20mccune.pdf>

Problem panels: Explore the various aspects of integrating trust in embedded systems from the perspective of the consumers, designers, integrators, and acquirers

Helen Gill, NSF
http://www.cert.org/tces/pdf/helen_gill.ppt.pdf

Ron Perez, AMD
<http://www.cert.org/tces/pdf/ron%20perez.pdf>

Duane Gilmour, AFRL
<http://www.cert.org/tces/pdf/duane%20gilmour.pdf>

Tom Schaffnit, Ford
<http://www.cert.org/tces/pdf/tom%20schaffnit.pdf>

Bob Thibadeau, Wave Systems
<http://www.cert.org/tces/pdf/robert%20thibadeau.pdf>

Breakout Sessions 1: Attack – Defend scenarios on example embedded systems

Wednesday, November 10, 2010

Focus is on Research

Reality Check on Embedded Systems constraints, Philip Koopman, CMU

<http://www.cert.org/tces/pdf/philip%20koopman%20part%202.pdf>

Report out on Current Research from researchers in academia, industry and government (part 1)

Tom Eisenbarth, Florida Atlantic

<http://www.cert.org/tces/pdf/thomas%20eisenbarth.pdf>

Kevin Fu, University of Massachusetts

<http://www.youtube.com/watch?v=shTj9WVhVyU>

Richard Han, University of Colorado

<http://www.cert.org/tces/pdf/richard%20han.pdf>

Trent Jaeger, Penn State University

<http://www.cert.org/tces/pdf/trent%20jaeger.pdf>

Ramesh Karri, Polytech, NYU

<http://www.cert.org/tces/pdf/ramesh%20karri.pdf>

Report out on Current Research from researchers in academia, industry and government (part 2)

Philip Koopman, CMU

<http://www.cert.org/tces/pdf/philip%20koopman.pdf>

Farinaz Koushanfar, Rice University

<http://www.cert.org/tces/pdf/farinaz%20koushanfar.pdf>

Insup Lee, University of Pennsylvania

<http://www.cert.org/tces/pdf/insup%20lee.pdf>

Roy Porter, AFRL

<http://www.cert.org/tces/pdf/roy%20porter.pdf>

Ed Suh, Cornell University

<http://www.cert.org/tces/pdf/ed%20suh.pdf>

Doug Blough, Georgia Tech

<http://www.cert.org/tces/pdf/doug%20blough.pdf>

Case Study in a directed attack – Stuxnet Dossier, Liam O Murchu, Symantec

<http://www.cert.org/tces/pdf/liam%20o%20murchu.pdf>

Breakout Session 2: Gap Analysis, Plenary

Appendix B – Attendees

Government

Bennett, Daren	NSA
Dean, Richard	DARPA
Gill, Helen	NSF
Gilmour, Duane	AFRL
Martin, William B.	NSA
May, Michael	OSD
Porter, Roy	AFRL

Industry

Aldridge, Hal	Sypris
Fansler, Aaron	Northrop Grumman
Forest, Tom	GM
Kalb, George	Northrop Grumman
Markham, Tom	Honeywell
McNamee, Dylan	Galois
Millen, Jon	Wave
O Murchu, Liam	Symantec
Peirce, Ken	GM
Perez, Ron	AMD
Prowell, Stacy	ORNL
Schaffnit, Tom	Ford
Thibadeau, Bob	Wave

Academia

Andrews, Archie	CMU
Blough, Doug	Georgia Tech
Chen, Peter	CMU
Eisenbarth, Thomas	Florida Atlantic
Farb, Michael	CMU
Fisher, David	CMU
Fu, Kevin	Univ. of Mass
Gligor, Virgil	CMU
Han, Richard	Univ. of Colorado
Jaeger, Trent	Penn State
Karri, Ramesh	NYU Poly
Koopman, Philip	CMU
Koushanfar, Farinaz	Rice University
Lee, Insup	Univ. of Pennsylvania
Lipson, Howard	CMU
McCune, Jonathan	CMU
Mok, Aloysius	Univ. of Texas-Austin
Perrig, Adrian	CMU
Shannon, Greg	CMU
Suh, G. Edward	Cornell
Szilagyi, Chris	CMU
Tsudik, Gene	UC Irvin

References

URLs are valid as of the publication date of this document.

[Agrawal 2007]

Dakshi Agrawal, et al. *Trojan Detection using IC Fingerprinting* (IBM Research Report RC24110[W0604-109]), April 21, 2006.

[Avizienis 2004]

Avizienis, Algirdas; Laprie, Jean-Claude; Randell, Brian & Landwehr, Carl. “Basic Concepts and Taxonomy of Dependable and Secure Computing.” *IEEE Transaction son Dependable and Secure Computing*, vol. 1, no. 1, Jan-Mar 2004.

REPORT DOCUMENTATION PAGE			<i>Form Approved</i> <i>OMB No. 0704-0188</i>	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave Blank)		2. REPORT DATE March 2011		3. REPORT TYPE AND DATES COVERED Final
4. TITLE AND SUBTITLE Trusted Computing in Embedded Systems Workshop			5. FUNDING NUMBERS FA8721-05-C-0003	
6. AUTHOR(S) Archie Andrews, Jonathan McCune				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Software Engineering Institute Carnegie Mellon University Pittsburgh, PA 15213			8. PERFORMING ORGANIZATION REPORT NUMBER CMU/SEI-2011-SR-002	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) HQ ESC/XPK 5 Eglin Street Hanscom AFB, MA 01731-2116			10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES				
12A DISTRIBUTION/AVAILABILITY STATEMENT Unclassified/Unlimited, DTIC, NTIS			12B DISTRIBUTION CODE	
13. ABSTRACT (MAXIMUM 200 WORDS) This report describes the November 2010 Trusted Computing in Embedded Systems Workshop held at Carnegie Mellon University. This workshop brought together various groups concerned with advancing research into improving the trustworthiness in embedded systems. The workshop format provided the opportunity to focus on embedded systems while examining the application of related trust technologies in order to foster collaborative approaches and information exchange in this area. Presentations and discussion addressed the capabilities and limitations of effectively employing trusted hardware-enabled components in embedded systems. This included, but was not restricted to, the following areas: new research and development in enabling trust in embedded systems, methods and techniques for establishing trust in embedded systems, lessons learned from research and development projects on embedded systems security, and gaps in current research. The workshop resulted in identification of gaps in current research and recommendations for potential research directions.				
14. SUBJECT TERMS trust, trust technology, trusted systems, embedded systems, trusted computing			15. NUMBER OF PAGES 32	
16. PRICE CODE				
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UL	